arXiv:1611.00090v1 [math.OC] 1 Nov 2016

On the Linear Convergence of the Cauchy Algorithm for a Class of Restricted Strongly Convex Functions

Hui Zhang*

November 2, 2016

Abstract

In this short note, we extend the linear convergence result of the Cauchy algorithm, derived recently by E. Klerk, F. Glineur, and A. Taylor, from the case of smooth strongly convex functions to the case of restricted strongly convex functions with certain form.

Keywords. Cauchy algorithm, gradient method, linear convergence, restricted strongly convex

AMS subject classifications. 90C25, 90C22, 90C20.

1 Introduction

The Cauchy algorithm, which was proposed by Augustin-Louis Cauchy in 1847, is also known as the gradient method with exact line search. Although this is the first taught algorithm during introductory courses on nonlinear optimization, the worst-case convergence rate question of this method was not precisely understood until very recently the authors of [1] settled it for strongly convex, continuously differentiable functions f with Lipschitz continuous gradient. This class of functions, denoted by $\mathcal{F}_{\mu,L}(\mathbb{R}^n)$, can be described by the following inequalities:

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|, \quad \forall x, y \in \mathbb{R}^n; \tag{1}$$

$$\langle \nabla f(x) - \nabla f(y), x - y\rangle \geq \mu\|x - y\|^2, \quad \forall x, y \in \mathbb{R}^n. \tag{2}$$

The gradient descent method with exact line search can be described as follows.

Algorithm 1 The gradient descent method with exact line search

Input: $f \in \mathcal{F}_{\mu,L}(\mathbb{R}^n)$, $x_0 \in \mathbb{R}^n$.
1: **for** $k = 0, 1, \cdots$, **do**
2: $\quad \gamma = \arg\min_{\gamma\in\mathbb{R}} f(x_i - \gamma\,\nabla f(x_i))$; // the exact linear search
3: $\quad x_{i+1} = x_i - \gamma\,\nabla f(x_i)$; // the gradient descent
4: **end for**

The authors of [1] obtained the following result, which settles the worst-case convergence rate question of the Cauchy algorithm on $\mathcal{F}_{\mu,L}(\mathbb{R}^n)$.

*College of Science, National University of Defense Technology, Changsha, Hunan, 410073, P.R.China. Corresponding author. Email: h.zhang1984@163.com

Theorem 1. *Let $f \in \mathcal{F}_{\mu,L}(\mathbb{R}^n)$, x_* a global minimizer of f on $\mathbb{R}^n$, and $f_* = f(x_*)$. Each iteration of the gradient method with exact line search satisfies*

$$f(x_{i+1}) - f_* \leq \left(\frac{L-\mu}{L+\mu}\right)^2 (f(x_i) - f_*), \quad i = 0, 1, \cdots . \tag{3}$$

For the case of quadratic functions in $\mathcal{F}_{\mu,L}(\mathbb{R}^n)$ with the form

$$f(x) = \frac{1}{2}x^T Q x + c^T x,$$

where $Q \in \mathbb{R}^{n\times n}$ is a positive definite matrix and $c \in \mathbb{R}^n$ is a vector, the result in Theorem 1 is well-known. The main contribution of [1] is extending the linear convergence of (3) from quadratic case to the case of $\mathcal{F}_{\mu,L}(\mathbb{R}^n)$. However, Theorem 1 can not completely answer what will happen when the matrix Q is positive semidefinite. In this note, we will further extend the linear convergence of (3) to a certain class of restricted strongly convex (RSC) functions, or more concretely to the functions that have the form of $f(x) = h(Ax)$, where $h \in \mathcal{F}_{\mu,L}(\mathbb{R}^m)$ and $A \in \mathbb{R}^{m\times n}$ with $m \leq n$, which cover the quadratic function $f(x) = \frac{1}{2}x^T Q x + c^T x$ with positive semidefinite matrix Q.

The concept of the restricted strongly convex was proposed in our previous paper [7, 8] and is strictly weaker than the strong convex; for detail and its recent development we refer the reader to [5, 3, 6]. It is not difficulty to see that $f(x) = h(Ax)$ is generally not strongly convex even $h(\cdot)$ is unless A has full column-rank. But it is always restricted strongly convex since it satisfies the following inequality that was proposed to define the RSC function:

$$\langle \nabla f(x), x - x' \rangle \geq \nu \|x - x'\|^2, \quad \forall x, y \in \mathbb{R}^n, \tag{4}$$

where x' is the projection point of x onto the minimizer set of $f(x)$, and ν is a positive constant. A decisive difference to strong convexity is that the set of minimizers of a RSC function need not be a singleton.

Now, we state our main result as follows:

Theorem 2. *Let $h \in \mathcal{F}_{\mu,L}(\mathbb{R}^m)$, $A \in \mathbb{R}^{m\times n}$ with $m \leq n$ and having full row-rank, $f(x) = h(Ax)$, x_* belong the minimizer set of f on $\mathbb{R}^n$, and $f_* = f(x_*)$. Denote $\kappa_h = \frac{\mu}{L}$ and $\kappa(A) = \frac{\lambda_{\min}(AA^T)}{\lambda_{\max}(AA^T)}$, where $\lambda_{\min}(AA^T)$ and $\lambda_{\max}(AA^T)$ stand for the smallest and largest eigenvalues of AA^T, respectively. Each iteration of the gradient method with exact line search satisfies*

$$f(x_{i+1}) - f_* \leq \left(\frac{2 - \kappa(A) - \kappa(A)\kappa_h}{2 - \kappa(A) + \kappa(A)\kappa_h}\right)^2 (f(x_i) - f_*), \quad i = 0, 1, \cdots . \tag{5}$$

When A is the identity matrix, Theorem 2 exactly recovers Theorem 1 since $\kappa(A) = 1$. From this sense, our main result is a further extension of Theorem 1. Recently, there appeared several papers [7, 3, 2, 6], the authors of which derived linear convergence results of gradient method with *fixed* step-length for RSC functions. The author of [3] analyzed the linear convergence rate of gradient method with exact linear search for RSC functions, but he only showed the existence of linear convergence rate. Our novelty here lies in the exact expression (5) for linear convergence.

For generally quadratic case, we have the following consequence.

Corollary 1. *Let $f(x) = \frac{1}{2}x^TQx + c^Tx$, where $Q \in \mathbb{R}^{n\times n}$ is a positive semidefinite matrix and $c \in \mathbb{R}^n$ is a vector; both of them are given. Assume that there is a vector x_* minimizing $f(x)$ on $\mathbb{R}^n$ and denote $f_* = f(x_*)$. Then, each iteration of the gradient method with exact line search satisfies*

$$f(x_{i+1}) - f_* \leq \left(1 - \frac{\lambda_{\min}^{++}(Q)}{\lambda_{\max}(Q)}\right)^2 (f(x_i) - f_*), \quad i = 0, 1, \cdots, \tag{6}$$

where $\lambda_{\min}^{++}(Q)$ stands for the smallest strictly positive eigenvalue of Q and $\lambda_{\max}(Q)$ is the largest eigenvalue of Q.

When Q is positive definite, applying Theorem 1 to the quadratic function $f(x) = \frac{1}{2}x^TQx+c^Tx$, we can get

$$f(x_{i+1}) - f_* \leq \left(\frac{\lambda_{\max}(Q) - \lambda_{\min}(Q)}{\lambda_{\max}(Q) + \lambda_{\min}(Q)}\right)^2 (f(x_i) - f_*), \quad i = 0, 1, \cdots.$$

But from Corollary 1, we can only obtain a worse linear convergence rate $\left(1 - \frac{\lambda_{\min}(Q)}{\lambda_{\max}(Q)}\right)^2$. Therefore, we wonder whether one can improve the rate in Corollary 1 from $\left(1 - \frac{\lambda_{\min}^{++}(Q)}{\lambda_{\max}(Q)}\right)^2$ to

$$\left(\frac{\lambda_{\max}(Q) - \lambda_{\min}^{++}(Q)}{\lambda_{\max}(Q) + \lambda_{\min}^{++}(Q)}\right)^2.$$

2 Proof

Proof of Theorem 2. We divide the proof into three steps.

Step 1. Denote $\lambda = \sqrt{\lambda_{\max}(AA^T)}$ and let $\widetilde{h}_\lambda(y) = h(\lambda y)$. Since $\nabla\widetilde{h}_\lambda(y) = \lambda \cdot \nabla h(\lambda y)$, from the definition of $h \in \mathcal{F}_{\mu,L}(\mathbb{R}^m)$, we have that

$$\|\nabla\widetilde{h}_\lambda(y) - \nabla\widetilde{h}_\lambda(z)\| \leq \lambda^2 L\|y - z\|, \quad \forall y, z \in \mathbb{R}^m; \tag{7}$$

$$\langle\nabla\widetilde{h}_\lambda(y) - \nabla\widetilde{h}_\lambda(z), y - z\rangle \geq \lambda^2\mu\|y - z\|^2, \quad \forall y, z \in \mathbb{R}^m. \tag{8}$$

Denote $\widetilde{u} = \lambda^2\mu.\widetilde{L} = \lambda^2 L$ and $\widetilde{A} = \lambda^{-1}A$; Then we can conclude that $\widetilde{h}_\lambda(y) \in \mathcal{F}_{\widetilde{\mu},\widetilde{L}}(\mathbb{R}^m)$ and

$$f(x) = h(Ax) = \widetilde{h}_\lambda(\widetilde{A}x).$$

Step 2. The iterates, generated by the gradient method with exact line search, satisfy the following two conditions for $i = 0, 1, \cdots$,

$$x_{i+1} - x_i + \gamma\nabla f(x_i) = 0, \quad \text{for some} \quad \gamma > 0; \tag{9}$$

$$\nabla f(x_{i+1})^T(x_{i+1} - x_i) = 0, \tag{10}$$

where the first condition follows from the gradient descent step, and the second condition is an alternative expression of the exact linear search step. Since $\gamma > 0$, it holds that successive gradients are orthogonal, i.e.,

$$\nabla f(x_{i+1})^T\nabla f(x_i) = 0, \ i = 0, 1, \cdots. \tag{11}$$

Note that $\nabla f(x) = \widetilde{A}^T \nabla \widetilde{h}_\lambda(\widetilde{A}x)$. Denote $y_i = \widetilde{A}x_i, i = 0, 1, \cdots$; Then, in light of the conditions (10) and (11), we can get that

$$\nabla \widetilde{h}_\lambda(y_{i+1})^T (y_{i+1} - y_i) = 0, \ i = 0, 1, \cdots ; \tag{12}$$

$$\nabla \widetilde{h}_\lambda(y_{i+1})^T \widetilde{A}\widetilde{A}^T \nabla \widetilde{h}_\lambda(y_i) = 0, \ i = 0, 1, \cdots . \tag{13}$$

Step 3. Following the arguments in [1], we consider only the first iterate, given by x_0 and x_1, as well as the minimizer y_* of $\widetilde{h}_\lambda(y) \in \mathcal{F}_{\widetilde{\mu},\widetilde{L}}(\mathbb{R}^m)$.

Set $h_i = \widetilde{h}_\lambda(y_i)$ and $g_i = \nabla \widetilde{h}_\lambda(y_i)$ for $i \in \{*, 0, 1\}$. Let $\epsilon = 1 - \kappa(A)$. Then, the following five inequalities are satisfied:

1: $h_0 \geq h_1 + g_1^T(y_0 - y_1) + \frac{1}{2(1-\widetilde{\mu}/\widetilde{L})} \left(\frac{1}{\widetilde{L}} \|g_0 - g_1\|^2 + \widetilde{\mu}\|y_0 - y_1\|^2 - 2\frac{\widetilde{\mu}}{\widetilde{L}}(g_1 - g_0)^T(y_1 - y_0) \right)$

2: $h_* \geq h_0 + g_0^T(y_* - y_0) + \frac{1}{2(1-\widetilde{\mu}/\widetilde{L})} \left(\frac{1}{\widetilde{L}} \|g_* - g_0\|^2 + \widetilde{\mu}\|y_* - y_0\|^2 - 2\frac{\widetilde{\mu}}{\widetilde{L}}(g_0 - g_*)^T(y_0 - y_*) \right)$

3: $h_* \geq h_1 + g_1^T(y_* - y_1) + \frac{1}{2(1-\widetilde{\mu}/\widetilde{L})} \left(\frac{1}{\widetilde{L}} \|g_* - g_1\|^2 + \widetilde{\mu}\|y_* - y_1\|^2 - 2\frac{\widetilde{\mu}}{\widetilde{L}}(g_1 - g_*)^T(y_1 - y_*) \right)$

4: $0 \geq g_1^T(y_1 - y_0)$

5: $0 \geq g_0^T g_1 - \epsilon \|g_0\| \|g_1\|$,

where the first three inequalities are the $\mathcal{F}_{\mu,L}$-interpolability conditions (see Theorem 4 in [4]), the fourth inequality is a relaxation of (12). It remains to show the fifth inequality. Indeed, in terms of (13), we have $g_o^T \widetilde{A}\widetilde{A}^T g_1 = 0$ and hence by the Cauchy-Schwartz inequality we can derive that

$$g_0^T g_1 = g_o^T (I - \widetilde{A}\widetilde{A}^T) g_1 \leq \|I - \widetilde{A}\widetilde{A}^T\| \cdot \|g_0\| \|g_1\|,$$

where I is the identity matrix. Since $\widetilde{A} = \lambda^{-1} A = \frac{A}{\sqrt{\lambda_{\max}(AA^T)}}$, it holds that

$$\|\widetilde{A}\widetilde{A}^T\| = \frac{\|AA^T\|}{\lambda_{\max}(AA^T)} \leq 1$$

and hence

$$\|I - \widetilde{A}\widetilde{A}^T\| = 1 - \lambda_{\min}(\widetilde{A}\widetilde{A}^T) = 1 - \frac{\lambda_{\min}(AA^T)}{\lambda_{\max}(AA^T)} = 1 - \kappa(A).$$

Therefore,

$$g_0^T g_1 \leq (1 - \kappa(A)) \|g_0\| \|g_1\| = \epsilon \|g_0\| \|g_1\|.$$

Since A is full row-rank, it is not difficulty to see that $0 \leq \epsilon < 1$. Now, we can repeat the arguments in the proof Theorem 5.1 in [1] and get

$$h_1 - h_* \leq \rho_\epsilon^2 (h_1 - h_*), \tag{14}$$

where $\rho_\epsilon = \frac{1-\kappa_\epsilon}{1+\kappa_\epsilon}$ and $\kappa_\epsilon = \frac{\widetilde{\mu}(1-\epsilon)}{\widetilde{L}(1+\epsilon)}$. After a simple calculus, we obtain

$$\rho_\epsilon = \frac{2 - \kappa(A) - \kappa(A)\kappa_h}{2 - \kappa(A) + \kappa(A)\kappa_h}.$$

Finally, noting that $f_* = h_*$ and for $i \in \{0, 1\}$ it holds

$$f(x_i) = h(Ax_i) = \widetilde{h}_\lambda(\widetilde{A}x_i) = \widetilde{h}_\lambda(y_i) = h_i,$$

we get

$$f(x_1) - f_* \leq \left(\frac{2 - \kappa(A) - \kappa(A)\kappa_h}{2 - \kappa(A) + \kappa(A)\kappa_h}\right)^2 (f(x_0) - f_*),$$

from which the conclusion follows. This completes the proof. □

Proof of Corollary 1. Since $\nabla f(x) = Qx + c$, by the assumption we have that $-Qx_* = c$. Let $m = \mathrm{rank}(Q)$ and let $Q = U\Sigma U^T$ be the reduced singular value decomposition, where $U \in \mathbb{R}^{n\times m}$ has orthogonal columns, and $\Sigma \in \mathbb{R}^{m\times m}$ is a diagonal matrix with the nonzero eigenvalues of Q as its diagonal entries. Denote $B = U\Sigma^{\frac{1}{2}}$, and $u = -\Sigma^{\frac{1}{2}}U^T x_*$; Then

$$f(x) = \frac{1}{2}x^T BB^T x + (B^T x)^T u = \frac{1}{2}\|B^T x + u\|^2 - \frac{1}{2}\|u\|^2.$$

Let $h(y) = \frac{1}{2}\|y + u\|^2 - \frac{1}{2}\|u\|^2$. Then we have $f(x) = h(B^T x)$ with $h(y) \in \mathcal{F}_{1,1}(\mathbb{R}^m)$ and B^T having full row-rank. Note that $\kappa_h = 1$ and

$$\kappa(B^T) = \frac{\lambda_{\min}(B^T B)}{\lambda_{\max}(B^T B)} = \frac{\lambda^{++}_{\min}(Q)}{\lambda_{\max}(Q)}.$$

Therefore, the desired result follows from Theorem 2. This completes the proof. □

Acknowledgements

This work is supported by the National Science Foundation of China (No.11501569 and No.61571008).

References

[1] E. d. Klerk, F. Glineur, and A. B. Taylor. On the worst-case complexity of the gradient method with exact line search for smooth strongly convex functions. *submitted to Optimization Letter, arXiv:1606.09365v1 [math.OC] 30 Jun 2016.*

[2] I. Necoara, Y. Nesterov, and F. Glineur. Linear convergence of first order methods for non-strongly convex optimization. *arXiv:1504.06298v2 [math.OC] 23 Apr 2015.*

[3] F. Schöpfer. Linear convergence of descent methods for the unconstrained minimization of restricted strongly convex functions. *SIAM J. Optim.*, 26(3):1883–1911, 2016.

[4] A. B. Taylor, J. M. Hendrickx, and F. Glineur. Smooth strongly convex interpolation and exact worst-case performance of first-order methods. *Math. Program., Ser. A*, DOI 10.1007/s10107-016-1009-3, 2016.

[5] H. Zhang. The restricted strong convexity revisited: analysis of equivalence to error bound and quadratic growth. *Optimization Letter, DOI:10.1007/s11590-016-1058-9, 2016.*

[6] H. Zhang. New analysis of linear convergence of gradient-type methods via unifying error bound conditions. *arXiv:1606.00269v4 [math.OC]*, 2016.

[7] H. Zhang and L. Cheng. Restricted strong convexity and its applications to convergence analysis of gradient-type methods. *Optimization Letter*, 9(5):961–979, 2015.

[8] H. Zhang and W. T. Yin. Gradient methods for convex minimization: better rates under weaker conditions. Technical report, CAM Report 13-17, UCLA, 2013.